Can ML predict airplane turbulence from wind data?

by Anna-Ida, Isabel & Rasmus

Context, motivation and goals

→ Why predict turbulence?

Passenger safety and comfort
Go-arounds and diversions are costly

→ at Nuuk airport:

Challenging terrain
Wind conditions complex (& dangerous)

→ Goal: **Predict experienced turbulence from wind data with machine learning**

https://www.sermitsiaq.ag/samfund/turbulens-ved-nuuk-kan-ende-i-flytragedie/521936

Topics during the Presentation

- Methods
  - Data
  - How to quantify turbulence?
  - Data Preparation
  - Machine Learning methods
- Results
- Conclusion
  - Summary
  - Outlook

AT / DK Buy subscription Log in

SERMITSIAQ

SOCIETY BUSINESS CULTURE SPORTS JOB

Two more Atlantic aircraft turn around - Kangerlussuaq is put into use

Air Greenland has fought a good fight to get all affected passengers flown across the Atlantic between Nuuk and Copenhagen in recent days. A flight from a Spanish airline flew part of the passengers to Nuuk over the weekend. On Monday, however, two more Atlantic aircraft had to turn around, and on Tuesday Kangerlussuaq will be temporarily put back into use for Atlantic flights.

https://www.sermitsiaq.ag/samfund/endnu-to-atlantfly-vender-om-kangerlussuaq-tages-i-brug/2179428

Data

- Flight recorder data
  - 3140 arrivals and 3039 departures from Nuuk airport in .csv format
  - 9 different types of planes, 2 weight classes
  - November 2024 - October 2025, sampled 8 times/second
  - Data on airplane position, speed, accelerations, ...
- Wind sensor (anemometer) data
  - 24 hours per day, 540 files in .csv format
  - sampled every 6 seconds
  - Data on wind speed, direction, ..., measured at the end of the runway

How to quantify turbulence?

- Turbulence ~ Amount of **vertical acceleration**
  - Perfectly smooth landing: Constant downward acceleration of 1g
  - For each landing, score the turbulence and use as label for the model, e.g.:

std turbulence = std(|vertical acceleration - 1g|)

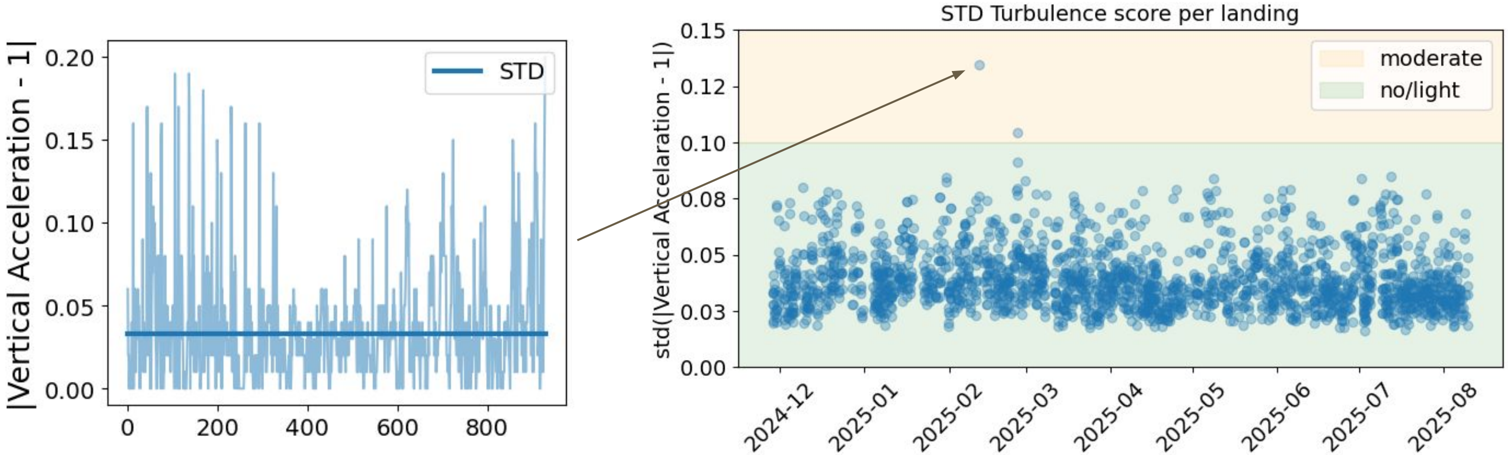

Features from wind data

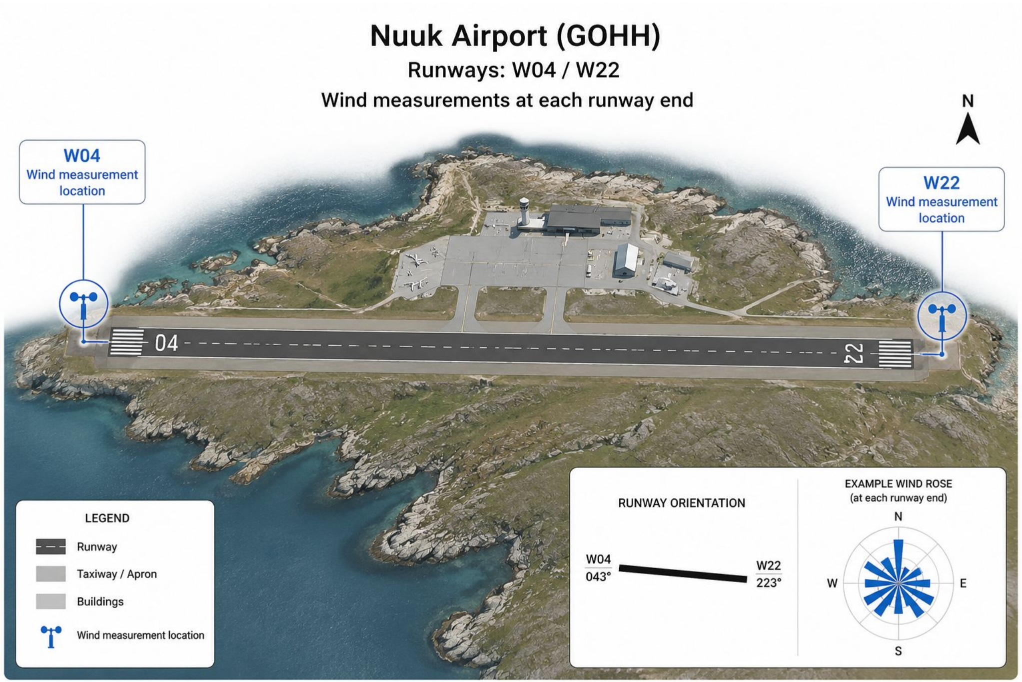

Define **wind features**, such as:

- Vector decomposition:

  u = speed * cos(angle)
  v = speed * sin(angle)
- Maximum, mean, range and deviation of wind speed
- Maximum, mean, range and deviation of wind direction
- Month, hour
- Approach direction
- …

Data Preparation

- Cleaning and streamlining, for example:
    - Different plane types have different file formats
    - Some fields had inconsistent formatting
    - Filter erroneous data points
- Only focus on 1500 ft (~ 460 meters) above the ground
- For each landing:
    - Calculate one **turbulence score**
    - Find **time window** for flights and extract corresponding **wind features**

Machine Learning Methods

- We tried: **XGBoost**, **LightGBM** and **LSTM**
- Workflow:
  - Optimization: Random Search or Optuna
  - Evaluation: RMSE and R^2 values (and visual inspection), SHAP values

Flight nr.	Wind speed	Wind speed std	other features	Turbulence score
0	1.785813	0.774385	xxx	0.05
…	…	…	…	…

XGBoost on Arrivals

- Model captures the **general trend** in turbulence, but
  - there seems to be only weak indicators for low turbulence events
  - underestimates sharp peaks, which are precisely the higher turbulence events we are trying to predict
  - model appears to underfit/regress towards the mean (dashed line)

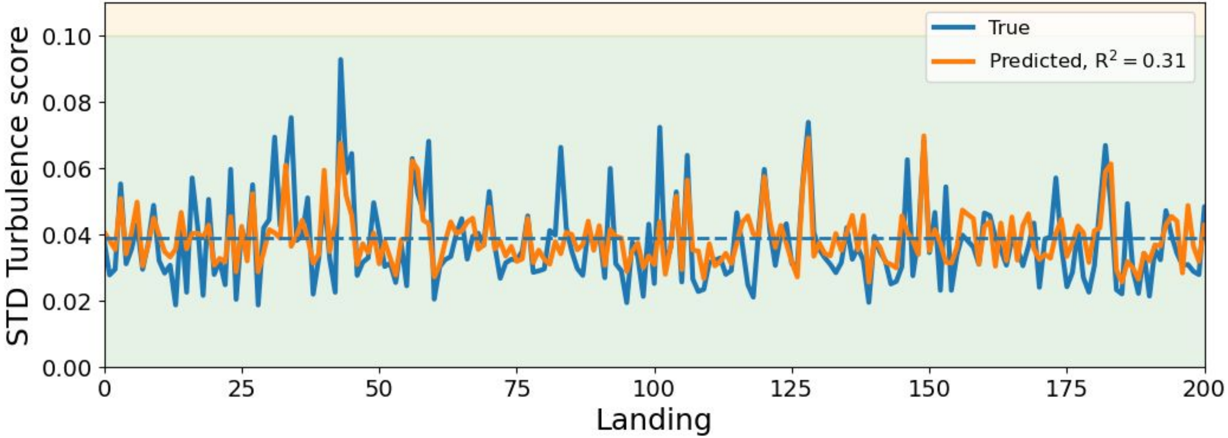

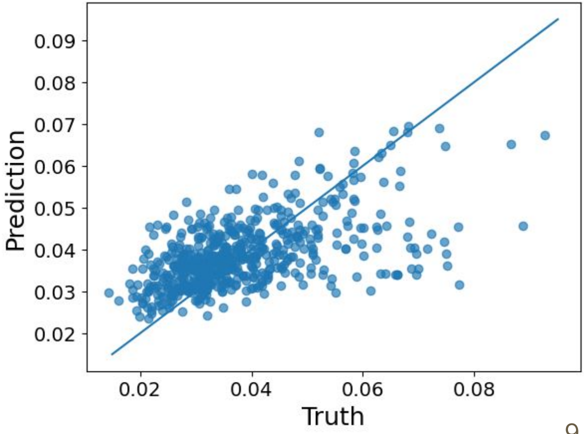

XGBoost on Arrivals

Permutation importance and SHAP values indicate that:

- Approach direction/**heading** matters a lot
- **Wind speed variability** and **change in wind direction** dominate
- Features calculated over **long time windows** are important
- **Seasonality** matters
- But: SHAP values are very small, no single feature is moving the prediction very much

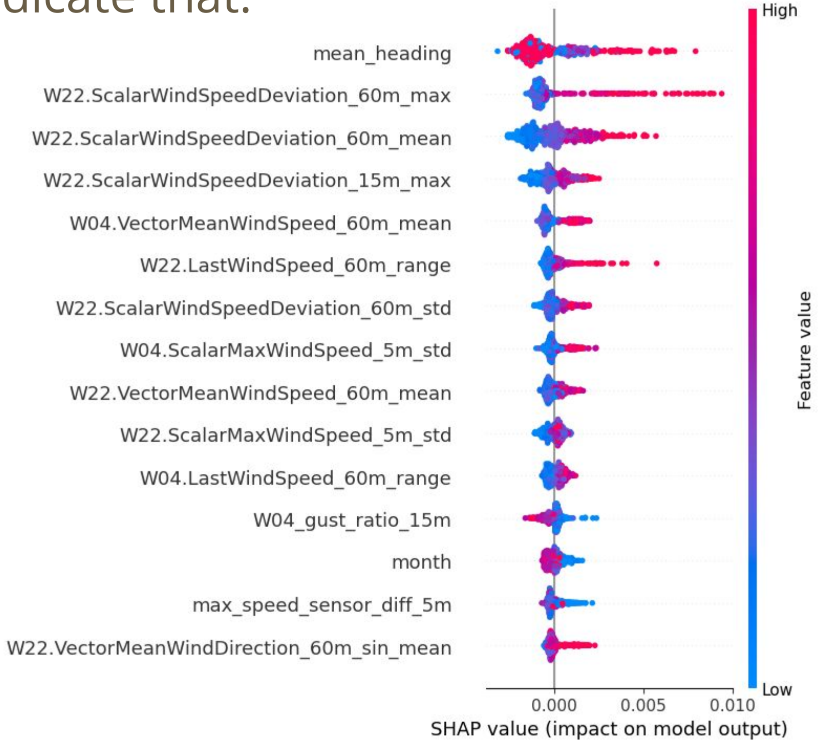

XGBoost hybrid model

- A bit better at capturing spikes

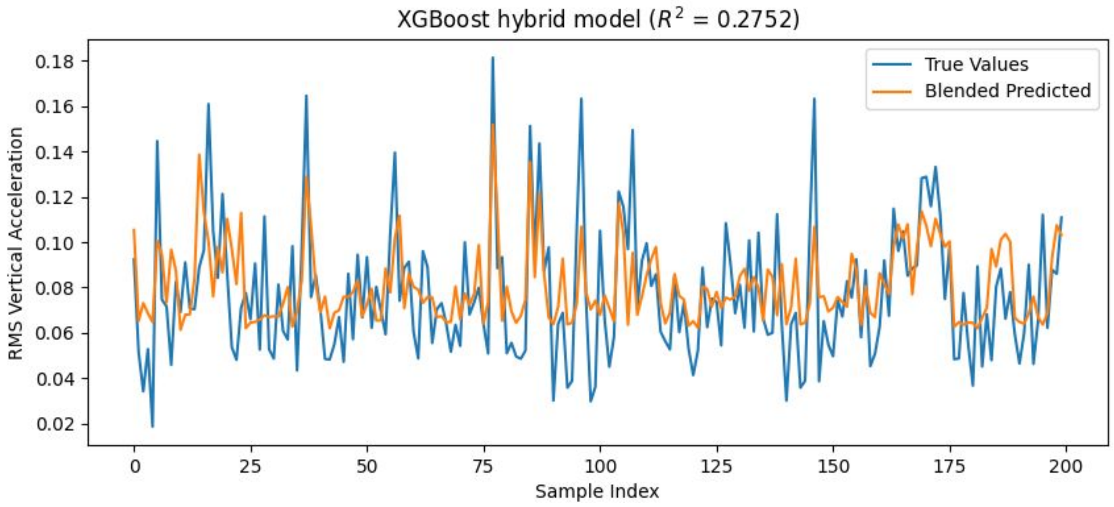

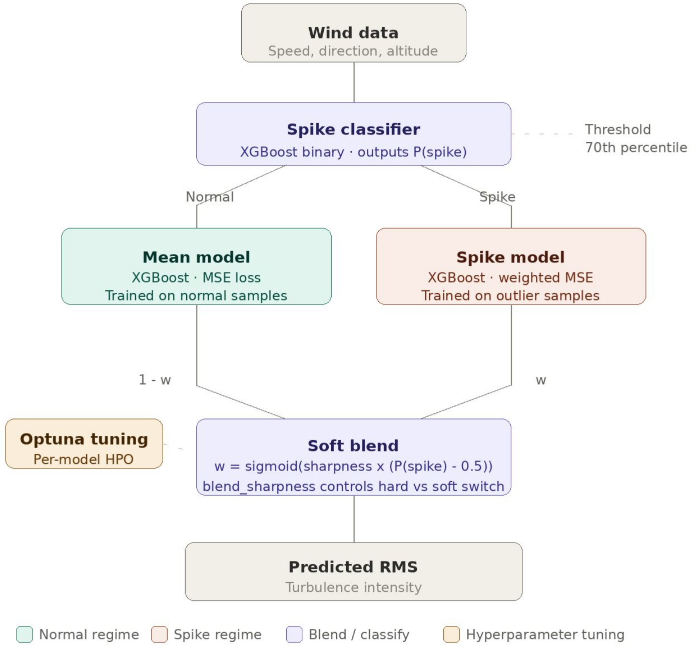

XGBoost hybrid model

- A bit better at capturing spikes
- ‘Floor’ where low true values do not lead to low predicted values

	Predicted moderate	Predicted No/Light
Moderate (n=31)	**8 (25.8%)**	**23 (74.2%)**
No/Light (n=249)	**7 (2.8%)**	**242 (97.2%)**

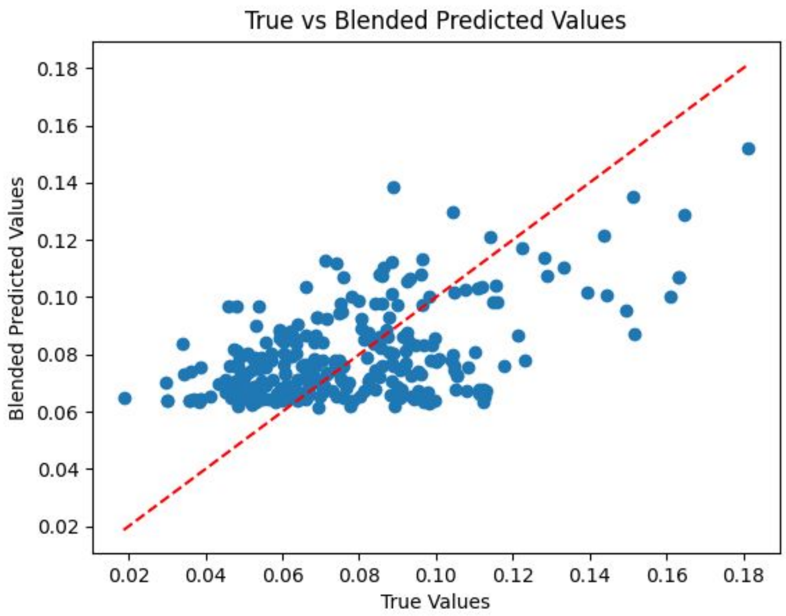

LSTM model

- 2793 arrivals, 120 time steps, 20 features → RMS turbulence output
- Bidirectional LSTM meant to capture temporal evolution of the wind

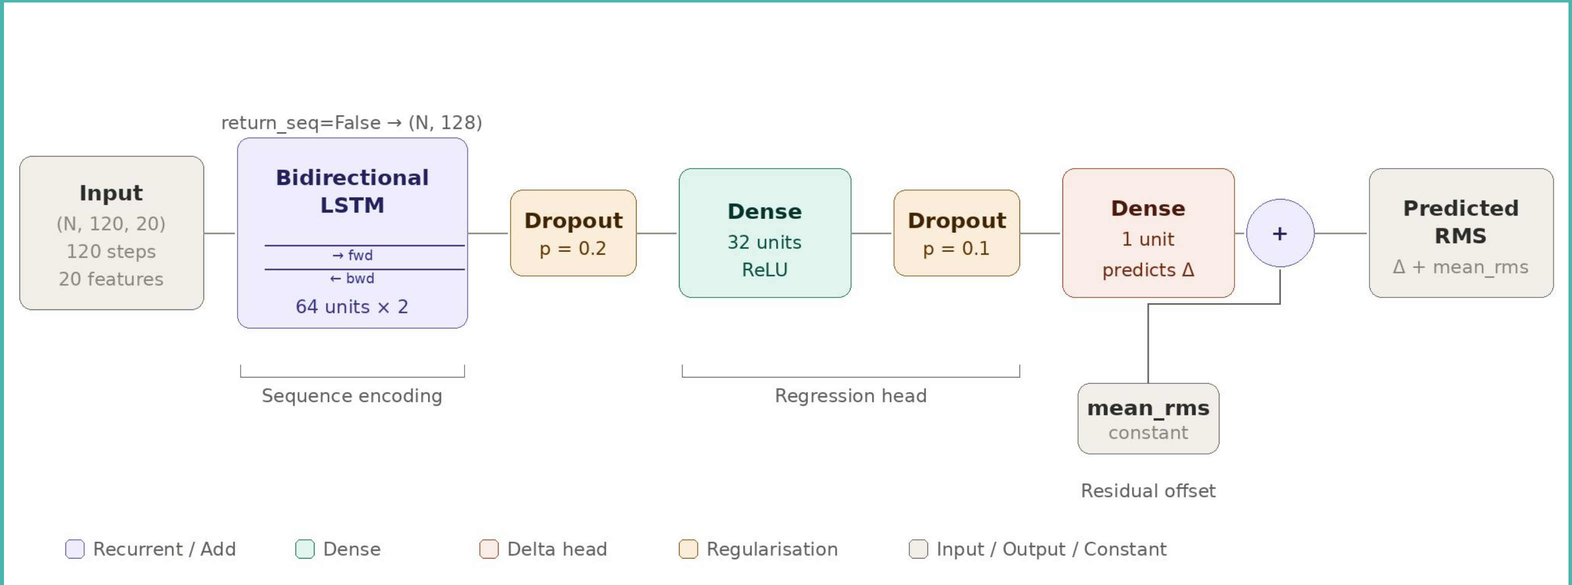

LSTM model

- 2793 arrivals, 120 time steps, 20 features → RMS turbulence output
- Bidirectional LSTM meant to capture temporal evolution of the wind
- Tends to underfit and regress heavily toward the mean
- Possibly needs far more data

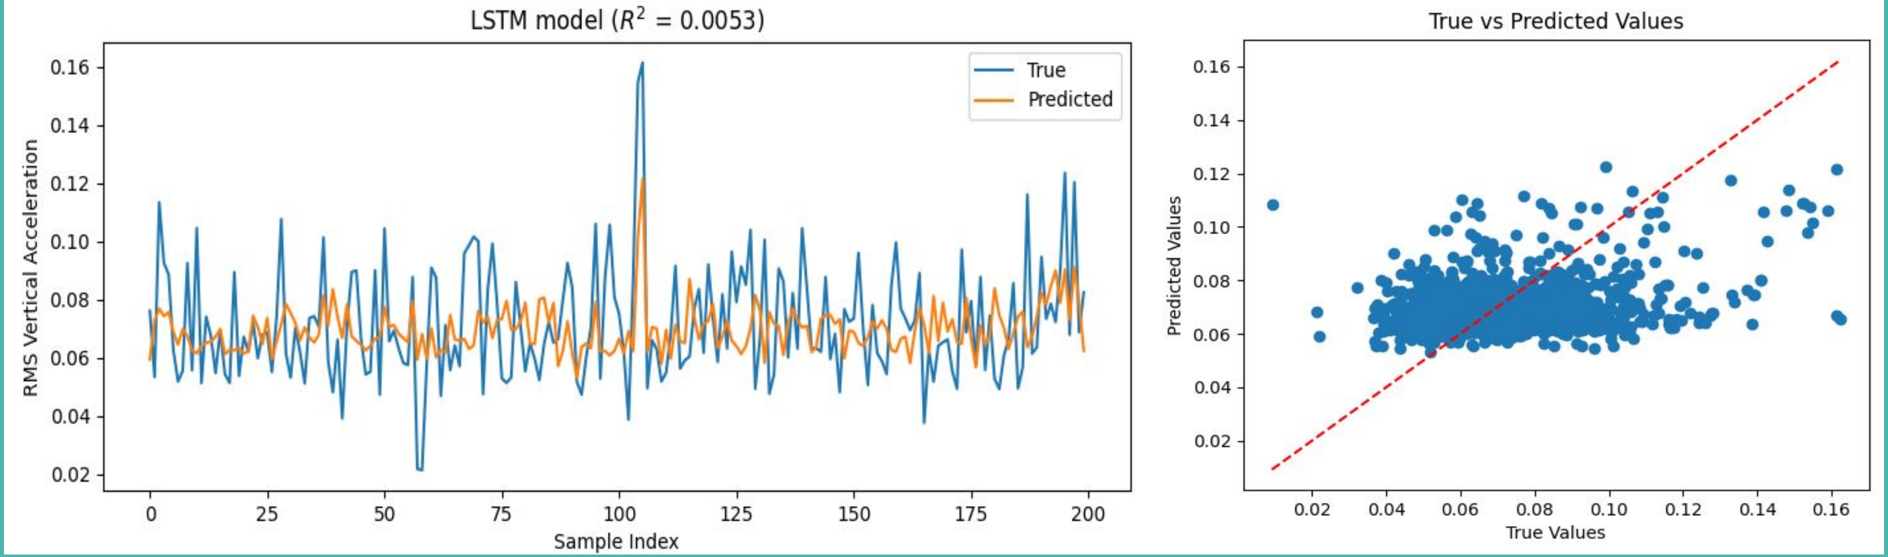

Results - Departures

for LightGB: RMSE: 0.009335340056159156 R2: 0.3471409453187534
for XGBoost: RMSE: 0.009189798547439837 R2: 0.36733890427970506

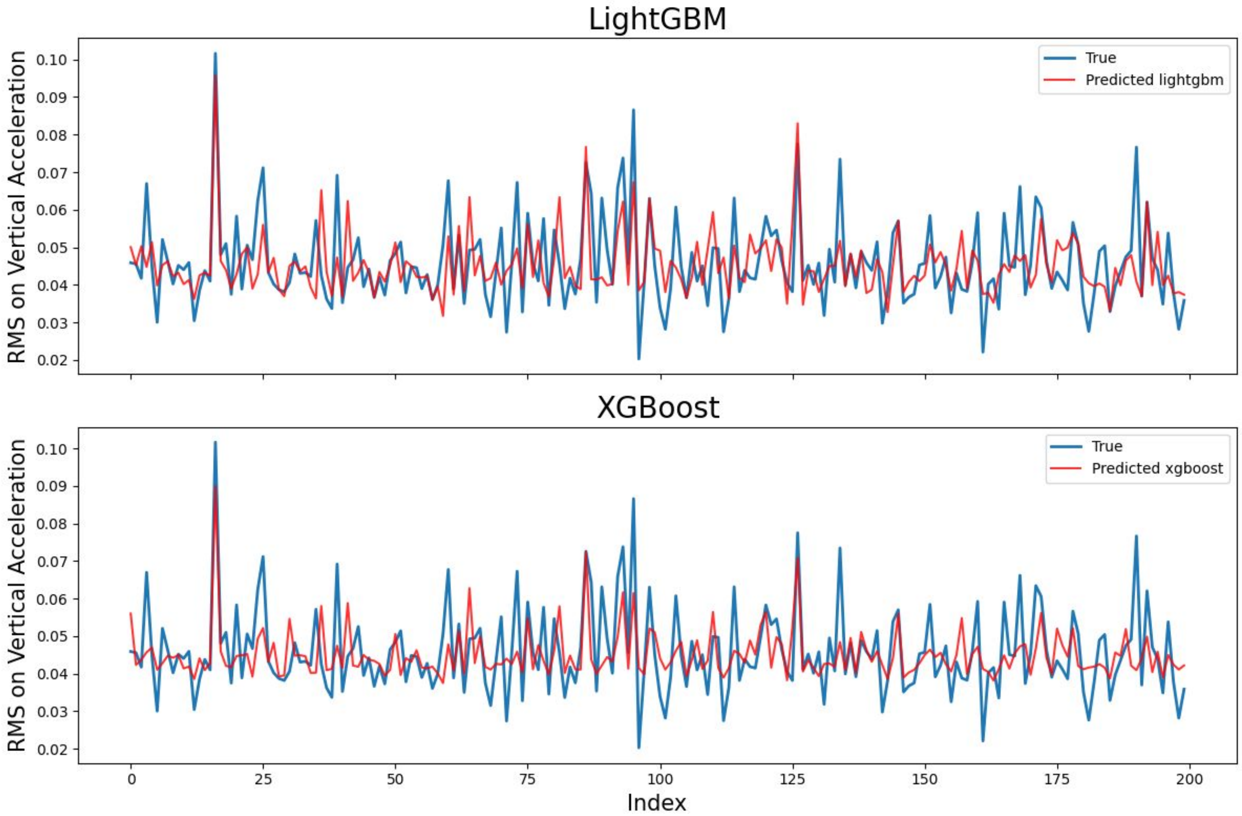

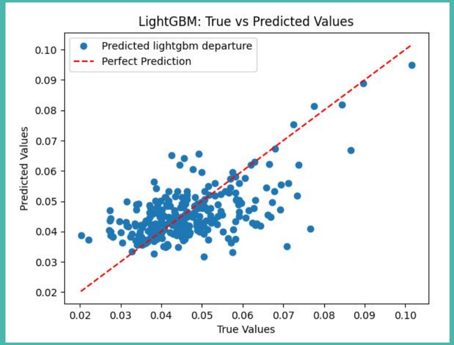

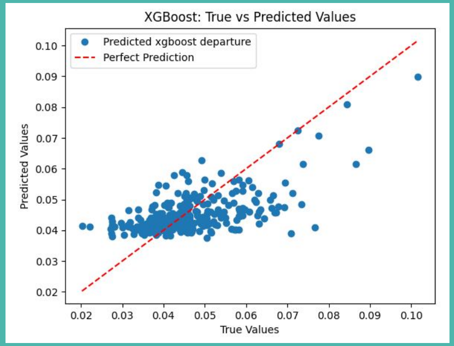

SHAP feature importance - Departure

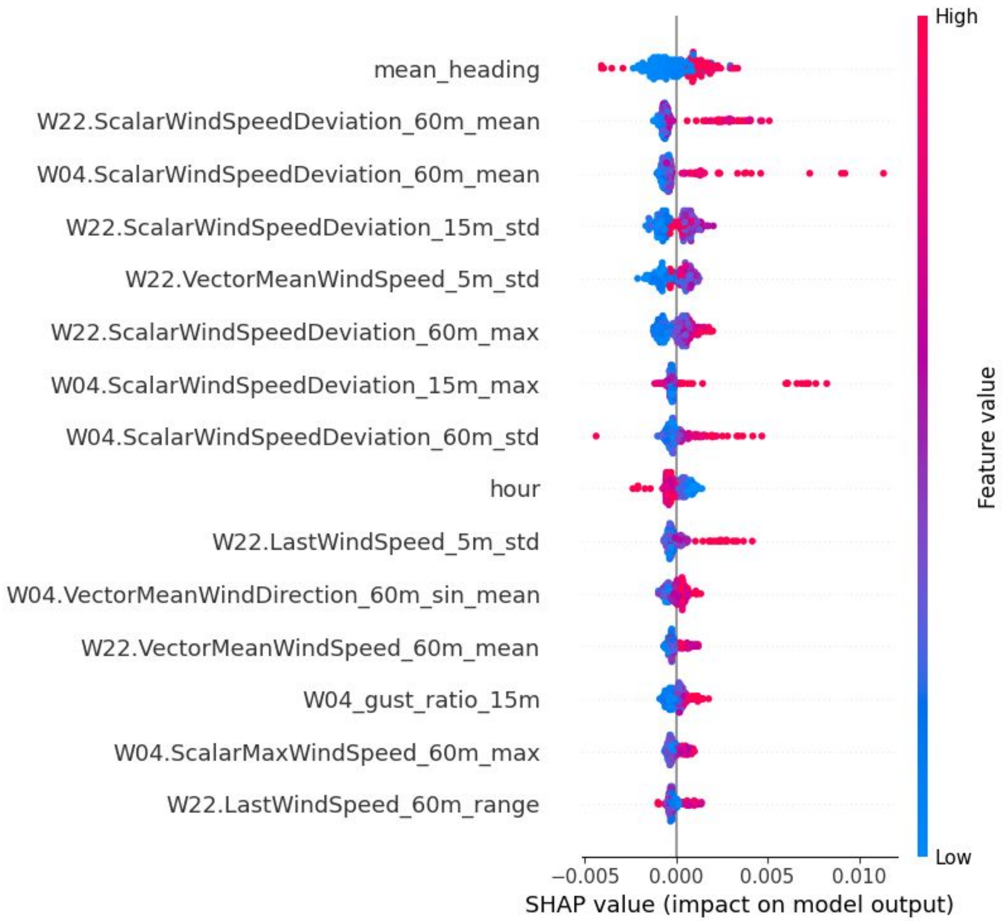

Challenges

- Complex, **large dataset** with some quality and formatting issues (and no documentation or domain knowledge)
- Is vertical acceleration really a **useful predictor for turbulence**?
  - RMS? Maximum? Standard deviation? Percentiles?
  - Just noise or is there actually a signal?
- Bias towards **favorable weather** and **low turbulence**
  - Our data contains only successful landings
  - Model never sees any medium to heavy turbulence
  - Limits how useful this approach is in practice

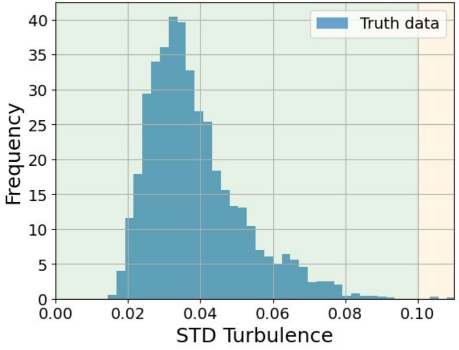

Summary

- Tried to predict turbulence experienced by airplanes in Nuuk from wind sensor data
- Model captures the turbulence level (R^2 = 0.3), but underpredicts the high-turbulence events that matter most for flight safety
- **Can ML help predict airplane turbulence from wind data?** Yes, but:
  - Calculating one turbulence score per landing is perhaps too simple
  - We may be missing important features
  - There is no moderate or severe turbulence data to train on

Outlook

- With more time, we would like to
  - Look at both arrival and departure data as a whole dataset
  - Look at the heavy planes, different plane types
  - Experiment with filtering the turbulence to get a clearer signal
- To really improve the prediction, we would need **more moderate and high turbulence data**.

... any Questions?

Appendix

Data filtering steps

1. Remove Radio_Altitude > 1500 ft
2. Filter out obviously wrong data (-> Plots)
   a. Arrivals: Some planes start on the ground and jump up (orange curve), filter this out
   b. Some planes with obviously wrong dates
3. Remove files that simply don't work
   a. Wrong data format (Especially for dates)
   b. Some plane types have the data formatted very differently, we skipped these for now to save time)

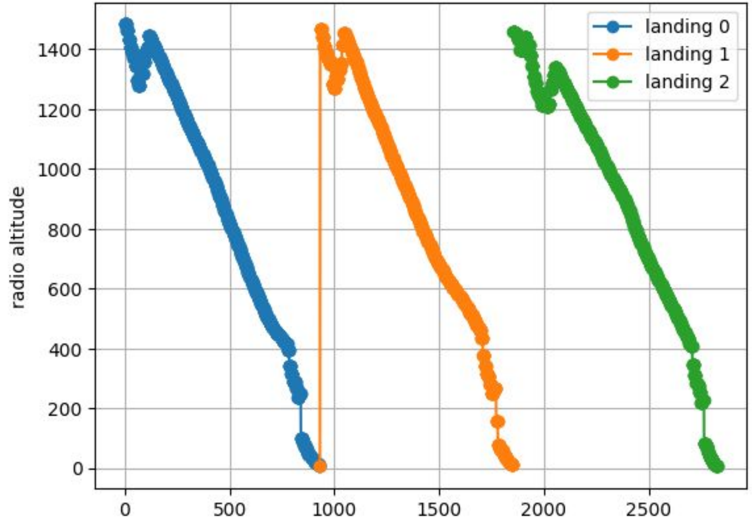

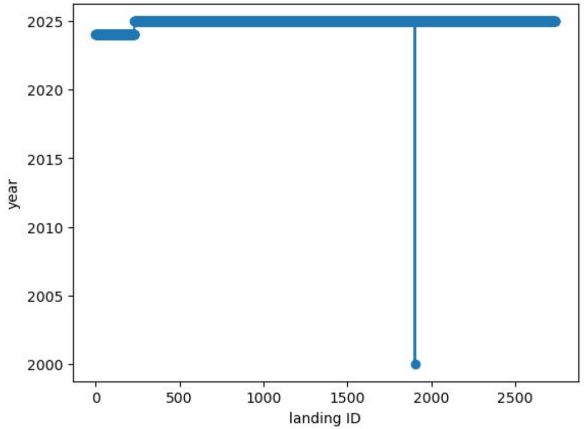

LightGBM on Arrival data

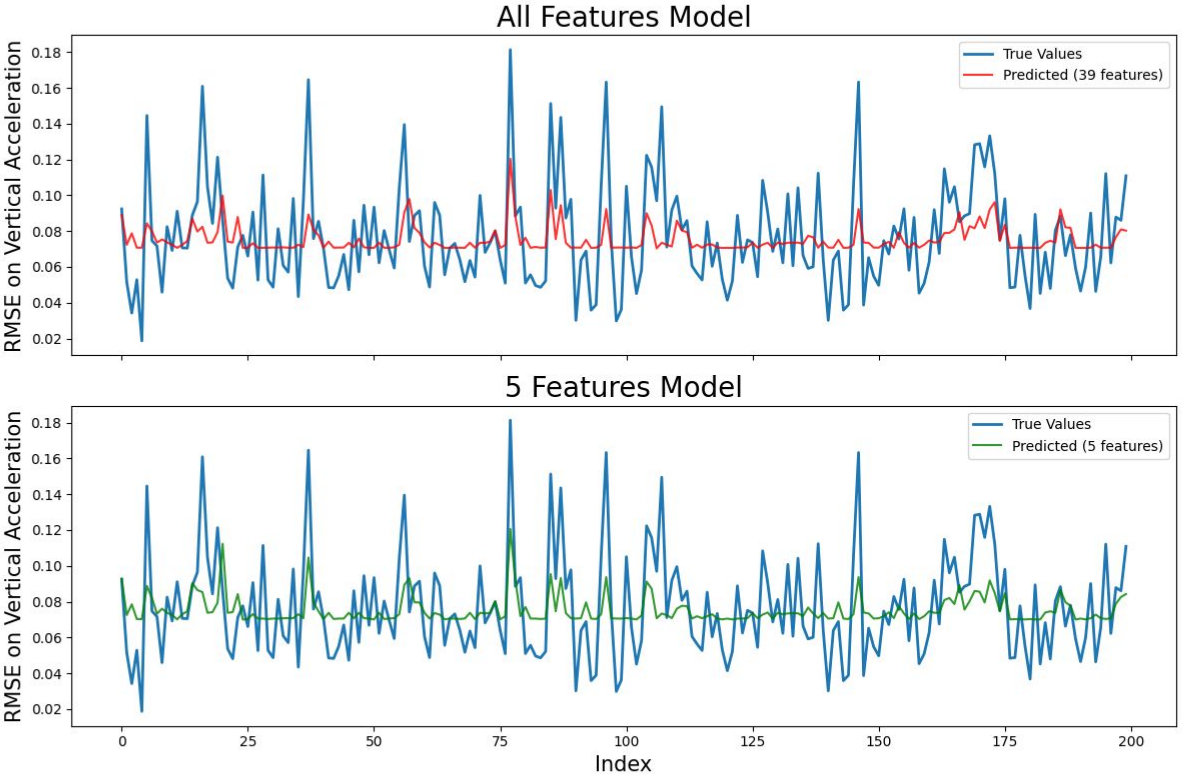

The five highest ranked features versus using all features available.

XGBoost on Arrivals

Top 15 features from **Permutation Importance (Random Forest)**:

14	W22.ScalarWindSpeedDeviation_60m_mean
97	mean heading
50	W22.ScalarWindSpeedDeviation_15m_max
16	W22.ScalarWindSpeedDeviation_60m_max
7	W04.ScalarWindSpeedDeviation_60m_mean
15	W22.ScalarWindSpeedDeviation_60m_std
0	W04.VectorMeanWindSpeed_60m_mean
19	W22.LastWindSpeed_60m_range
53	W22.LastWindSpeed_15m_std
12	W22.VectorMeanWindSpeed_60m_mean
35	W22.VectorMeanWindDirection_60m_sin_mean
11	W04.LastWindSpeed_60m_range
67	W04.VectorMeanWindDirection_15m_sin_mean
95	month
79	W22.ScalarWindDirectionDeviation_5m_std

Top 15 features from **XGBoost** (after optimization with Random Search):

	feature	importance
3	W22.ScalarWindSpeedDeviation_60m_max	0.050060
7	W22.LastWindSpeed_60m_range	0.035421
24	W04.LastWindSpeed_15m_range	0.024991
58	W04.ScalarMaxWindSpeed_15m_std	0.019002
0	W22.ScalarWindSpeedDeviation_60m_mean	0.017770
34	W04.LastWindSpeed_60m_std	0.017436
42	W04.LastWindSpeed_15m_std	0.017008
6	W04.VectorMeanWindSpeed_60m_mean	0.015635
68	W04.ScalarWindSpeedDeviation_15m_max	0.014304
90	W04.ScalarMeanWindDirection_15m_max_change	0.014229
2	W22.ScalarWindSpeedDeviation_15m_max	0.014080
4	W04.ScalarWindSpeedDeviation_60m_mean	0.013261
9	W22.VectorMeanWindSpeed_60m_mean	0.012885
5	W22.ScalarWindSpeedDeviation_60m_std	0.012802
95	max_speed_sensor_diff_5m	0.012101

XGBoost on Arrivals

Best parameters for XGBoost, found via **Random Search**:

subsample=0.7
n_estimators=300
max_depth=7
learning_rate=0.05
colsample_bytree=0.7

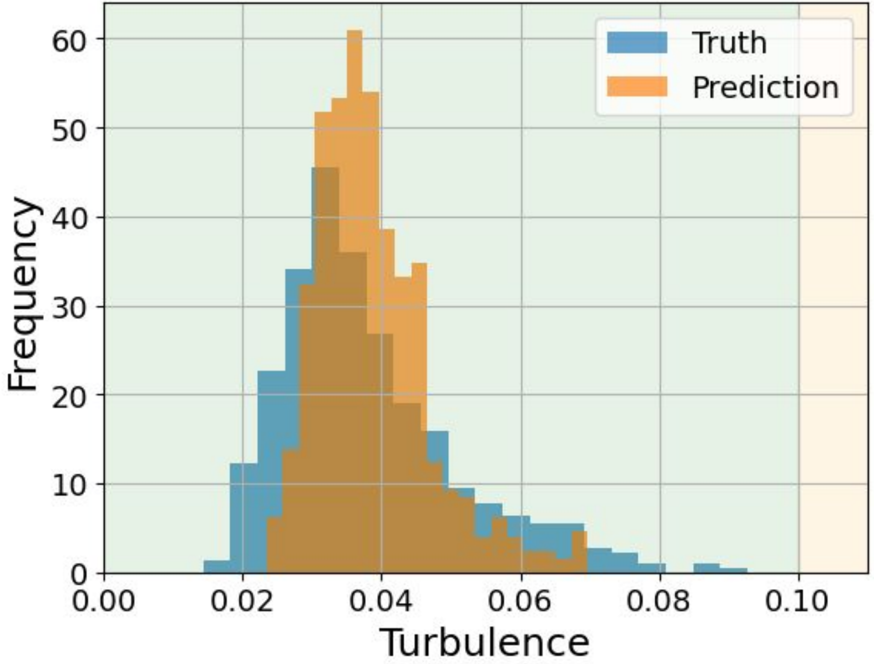

XGBoost Hybrid Model hyperparameters

CLASSIFIER: {‘n_estimators’: 5000, 'max_depth': 4, 'learning_rate': 0.034575879047178036, 'subsample': 0.6897782482740652, 'colsample_bytree': 0.6926504876850375}

MEAN REGRESSOR: {‘n_estimators’: 5000, 'max_depth': 8, 'learning_rate': 0.011406317887383348, 'subsample': 0.603904781824156, 'colsample_bytree': 0.8510582023021905}

SPIKE REGRESSOR: {‘n_estimators’: 5000, 'max_depth': 7, 'learning_rate': 0.037336170270069094, 'subsample': 0.6400568787844034, 'colsample_bytree': 0.7015833093192322}

Current calculation of turbulence

Two used features:

- Looking at speed of wind and its direction.
- Using pilot reports for comparison of threshold.